

# Capacity Performance of Dynamic Provisioning in Optical Networks

Ramu Ramamurthy, Zbigniew Bogdanowicz, Shahrokh Samieian, Debanjan Saha, Bala Rajagopalan, Sudipta Sengupta, Sid Chaudhuri, *Member, IEEE*, and Krishna Bala

**Abstract**—This paper describes an architecture and analyzes the performance of dynamic provisioning of lightpaths in an optical network. In dynamic provisioning, a lightpath is set up in real-time without rearranging the working and protection routes of existing lightpaths, and without the knowledge of future lightpath provisioning events. This paper develops a general model of the physical topology of the optical network, and outlines routing approaches for dynamic provisioning of lightpaths. It analyzes via simulations the performance of dynamically provisioned unprotected, 1 + 1 protected and mesh-restored lightpaths. The analysis of the efficiency of network utilization of dynamic provisioning focuses on the spare capacity needed for protection, and in particular focuses on the impact of sharing of wavelength channels for mesh-restored lightpaths. The main conclusion from the performance studies is that significant capacity gains are achieved with sharing of wavelength-channels for mesh-restored lightpaths with dynamic provisioning even for sparse topologies, and even at moderate loads.

**Index Terms**—Communication system performance, internet-working, optical communication.

## I. INTRODUCTION

TO accommodate the exponential growth of the Internet, transport networks based on wavelength division multiplexing (WDM) technology [1] are increasingly being deployed in carrier networks. WDM technology harnesses the enormous bandwidth of the fiber (potentially tens of Terabits per second) by multiplexing hundreds of optical channels each of which operate at electronic speeds (several Gigabits per second). Point-to-point WDM links that multiplex several tens of optical channels are currently deployed in the carrier networks.

An optical network comprises of optical switches interconnected in a general mesh network configuration by point-to-point WDM links. An optical switch can switch an entire wavelength channel (e.g., at OC-48, OC-192 rates) from input ports to output ports. In general, optical switches can be purely optical, or electronic or a combination of optical and electronic depending on the degree to which signals remain in the optical or electronic domains within the switch. In this paper it is assumed that the optical switches allow for full wavelength conversion, i.e., any wavelength channel on an input port can be switched to any wavelength channels on output ports.

The basic service provided by an optical network is to provision wavelength circuits called “lightpaths” between the client Network Elements (NEs). A client NE is a device such as an internet protocol (IP) router, an ATM switch, or an add-drop multiplexer that connects to the optical switches via drop ports of the optical switch. A lightpath is an optical channel from a drop port at an optical switch to a drop port at a remote optical switch that traverses multiple intermediate optical switches. An example of a lightpath is an OC-48 circuit setup between two client IP routers connected to the optical network. An optical network with optical switches allows for dynamically setting up and tearing down lightpaths. This enables a network operator to provision a lightpath in a matter of seconds instead of days or months. A lightpath provisioning request may be initiated by the network operator from a management system. Alternatively, a lightpath provisioning request may be initiated by a client NE from the User Network Interface (UNI). A lightpath is provisioned by first finding a route for the lightpath in the optical network with available capacity, and then configuring the optical switches along the lightpath route.

There are different kinds of failures which can disrupt lightpath services. The most common failures are single channel failures caused by the failures of equipment at a port of an optical switch. Fiber cuts can disrupt all (potentially hundreds) the optical channels that are carried on the optical fiber. An optical switch is designed to have redundant equipment (including redundant switch matrices), and are resilient to equipment failures. Switch failures though rare can cause the failure of all channels that are adjacent to the optical switch.

This paper considers the following three types of lightpaths:

- 1) *unprotected* lightpaths;
- 2) *1 + 1 protected* lightpaths;
- 3) *mesh-restored* lightpaths.

An unprotected lightpath is not protected upon the failure of any equipment (fiber-links, transceivers, etc.) along the lightpath route. A 1 + 1 protected lightpath has a working route and a diversely routed protection route. The wavelength channels for the working and diversely routed backup route are dedicated for a 1 + 1 protected lightpath. The working and protection routes of a 1 + 1 protected lightpath can be

- 1) *edge-disjoint* which allows the lightpath to recover from single-link failures or
- 2) *node-disjoint* which allows the lightpath to recover from single-link and single-node failures.

Manuscript received June 7, 2000; revised October 2, 2000.

The authors are with Tellium, Inc., Oceanport, NJ 07757 USA (e-mail: ramu@tellium.com; zbogdanowicz@tellium.com; ssamieian@tellium.com; dsaha@tellium.com; braja@tellium.com; sudipta@tellium.com; sc@tellium.com; kbala@tellium.com).

Publisher Item Identifier S 0733-8724(01)00407-8.

A mesh-restored<sup>1</sup> lightpath has a working route and a diversely routed backup route. The wavelength channels on the working route of the mesh-restored lightpath are dedicated for that lightpath and carry user traffic under normal operating conditions. The wavelength channels on the backup route for the mesh restored lightpath are shared among different mesh-restored lightpaths. Wavelength channels are shared in to ensure that any single fiber-link failure on the working route of any mesh-restored lightpath can be restored.

The objective of this paper is to present an architecture for dynamic provisioning, and examine its capacity utilization efficiency. In this model of dynamic provisioning, it is assumed that each arriving lightpath is provisioned in real-time subject to the following constraints:

- 1) existing (working and backup) lightpaths cannot be rerouted;
- 2) future lightpath requests are not known.

The dynamic provisioning model corresponds to a network operation environment where the time-scale for lightpath provisioning is of the order of seconds. This contrasts with the static provisioning model where a set of lightpath demands are known ahead of time for a future period of time; then, the entire set of lightpaths can be provisioned in a nonreal-time manner, where, there are no time-constraints for provisioning each lightpath. The dynamic provisioning model enables an environment where client devices of the optical network such as IP routers can dynamically request lightpaths from the optical network. The dynamic provisioning model also allows network operators to rapidly deploy lightpaths, without disrupting existing lightpath services. We assume that over longer time periods, the network operator may reoptimize the lightpath routes in an offline manner by globally rerouting the backup paths (or both the working and backup) of a set of existing lightpaths.

The outline for the paper follows. Section II describes the network architecture, and defines the mechanisms and protocols used for discovering the topology information, and dynamically provisioning a lightpath. Section III describes the modeling of the topology information, and the approach for routing lightpaths. Section IV presents the results of the performance studies. Section V concludes the paper.

## II. NETWORK ARCHITECTURE

### A. Physical Architecture

Fig. 1 illustrates the architecture of an optical network where multiple optical switches are interconnected via WDM links in a general mesh interconnection pattern. An optical switch has multiple ports and is capable of switching an optical channel from an input port to an output port. There is a distinction made between drop ports and network ports, although they represent the same physical circuit card. Access devices such as IP routers, ATM switches, or add-drop multiplexers are connected to the drop ports of an optical switch using standard interfaces. Access devices are connected to their peers using dynamically es-




Fig. 1. Optical network.




Fig. 2. Optical network physical topology with SRGs.

tablished lightpaths across the optical network. The interaction between access devices and optical switches for dynamic provisioning of lightpaths occurs with well-defined routing and signaling interfaces that comprise the UNI. Neighboring optical switches may be interconnected by multiple optical channels. The ports used to interconnect neighboring optical switches are denoted as network ports.

### B. Shared Risk Groups (SRGs)

A set of optical channels between neighboring optical switches that have the same level of risk of failure is called a shared risk group (SRG) [2]. For example, all the channels that are multiplexed onto a WDM fiber-link is a SRG since the failure of the fiber-link simultaneously affects all the channels that are carried over that fiber. A set of optical channels between neighboring optical switches can, in general, belong to multiple SRGs. For example, Fig. 2 illustrates a physical topology with four SRGs. SRG-1, SRG-2, and SRG-3 refer to the conduits that emerge from optical switches A, B, and C. SRG-4 refers to the WDM fiber link. The concept of SRGs allows for the definition of diverse routing in an optical network. For example, two lightpaths are diversely routed if their routes do not contain any common SRGs. This definition of diverse routing implies that any single SRG failure on one lightpath does not affect the diversely routed lightpath. If two lightpaths  $p_1$  and  $p_2$  do not have any SRGs in common, they are SRG-disjoint. In this paper, it is assumed that the fundamental unit of failure protection is a single SRG failure, i.e., for mesh-restored lightpaths, the failure of any single SRG on its working routes

<sup>1</sup>The term “mesh-restoration” used in this paper is also called shared-restoration in the literature.

will be restored. This guarantee can be met by ensuring that the working and backup routes for mesh-restored lightpaths are SRG-disjoint.

SRGs are configured by the network operator with the knowledge of the physical fiber plant of the optical network. In general, SRGs can be hierarchical. For example, the optical channels within a fiber, the fibers within a cable, and the cables within a conduit form a three level SRG hierarchy. In this paper it is assumed that by default, all channels between adjacent optical switches belong to a single SRG. In this default case, an SRG is synonymous with the link comprising all the channels between neighboring optical switches. When SRGs take on default values, fundamental unit of failure protection is against a single-link failure, and in this case mesh-restored lightpaths protect against all single-link failures in the optical network.

### C. Resource Discovery and Maintenance

The task of the topology discovery consists of two main components. First, neighboring switches exchange local connectivity information with each other via a neighbor discovery protocol. Then, an enhanced version of a link-state protocol such as open shortest path first (OSPF) [3] is used to disseminate local connectivity and network resource availability information throughout the network. Each switch in the network can create a network topology database of its own by piecing together the connectivity information advertised by the switches participating in the OSPF protocol.

1) *Neighbor Discovery Protocol*: Neighbor Discovery Protocol (NDP) [4] runs between two optical switches that peer at the link-layer and uses link-layer overhead bytes to exchange port and channel-state information. NDP runs on all network ports unless it is explicitly disabled. The primary purpose of NDP is to discover the switch ID of the peer optical switch and the port ID of the port on the remote end of the optical channel. Additionally, it also exchanges other optional configuration information, such as SRGs that an optical channel belongs to. The information gathered by NDP is reflected in the port-state database maintained by each switch. At the time of initialization the port-state entry corresponding to a port at an optical switch consists of the following information:

- local attributes (switch ID, port ID, port type): local attributes are autoconfigured;
- SRGs of the channels on the port are configured manually. An optical channel (port) may belong to multiple SRGs;
- remote port ID and remote switch ID are learned from NDP.

The port-state database from the NDP protocol supplements the link-state database from the OSPF protocol to provide sufficient information needed for the network topology discovery and lightpath management.

2) *Optical Extensions to OSPF*: An enhanced version of OSPF is used for topology discovery. Each optical switch is identified by a single IP address, which is equivalent to the router ID in the OSPF terminology. All switches are assumed to be within the same administrative domain and a single area. It is assumed that all the optical channels between two switches are abstracted

as a link bundle and appear as a single unnumbered point-to-point segment (link in OSPF terminology) to OSPF. It is also assumed that there exists a single IP control channel between two neighboring switches. The control channel is implemented using the link-layer overhead bytes in the optical channels. As long as at least one of the optical channels is operational between two switches, the IP control channel is assumed to be operational. The OSPF link-state database in its native form only provides connectivity information among NEs. Dynamic lightpath provisioning needs a more detailed link-state database that includes the number and the type of optical channels available between neighboring optical switches. Additionally, the link state database needs to capture SRG information to enable diverse-routing of lightpaths. The OSPF link-state database also has to capture the resource class an optical channel belongs to. For example, channels dedicated to working paths may not be shared between two paths while channels dedicated to shared-backup paths may be shared among multiple paths. Consequently, when a free channel is allocated to a working path it is no longer available and OSPF update should capture this change in state. On the other hand, when a free channel is allocated to a shared-backup path, it can no longer be assigned to a working path, but it can be assigned to a shared backup path and should be reflected in the OSPF link state update. OSPF opaque Link State Advertisement (LSA) provides a generalized mechanism to carry additional information. Opaque LSAs are used to disseminate optical resource related information. Standard link-state database flooding mechanisms are used for distribution of these opaque LSAs.

The topology database maintained at the NE consists of a set of opaque LSAs which describe the summarized topology of the optical network for use by the routing algorithms. The topology database consists of one entry for each channel group that is adjacent to a switch. The entry for a Channel Group that is broadcast in an optical LSA by a switch represents one directional connectivity between the local switch and a remote switch and contains the following fields:

- SRG IDs to which the Channel Group belongs;
- Local Switch ID;
- Remote Switch ID;
- for each channel rate (OC-48, OC-192);
  - number of available channels;
- cost assigned to Channel Group;
- propagation delay on a channel in the channel group (used by the routing algorithm to satisfy propagation delay constraints, and restoration time constraints for lightpaths).

Fig. 3 illustrates the topology database for the optical network topology depicted in Fig. 2.

The quality of a lightpath route in the optical network is determined by a cost model. A cost model assigns costs to channels in the network that represents the cost of using the channel in the route. The quality of the route is the sum of costs of all channels in the route. A user defined cost may be assigned to channels by the network operator to represent the fiber-mileage, fiber leasing cost, or some other metrics.

### D. Lightpath Attributes

A lightpath is characterized by the following attributes.

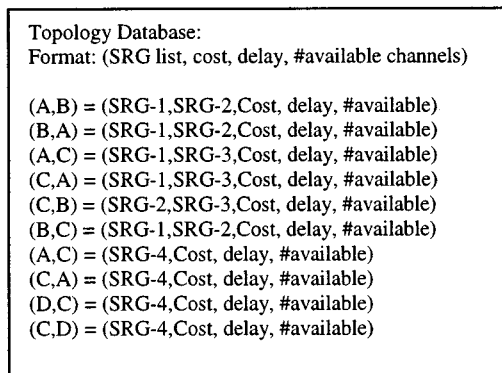


Fig. 3. Topology database consisting of a set of opaque LSAs.

- Lightpath ID: An identifier for the lightpath that uniquely identifies the lightpath from all other lightpaths in the optical network.
- Ingress and egress lightpath termination points: The ingress termination point is identified the ingress switch id, the ingress port ID, and ingress channel ID. Similarly, the egress termination point is identified by the egress switch ID, egress port ID, and the egress channel ID.
- Channel rate: e.g., OC-48, OC-192.
- Restoration attributes: The following restoration attributes are defined:
  - 1) unprotected;
  - 2) 1 + 1 protected (node-diverse, and SRG-diverse);
  - 3) mesh-restored.

Unprotected lightpaths are not restorable upon failures. 1 + 1 protected lightpaths are protected using diversely routed node (or SRG)-disjoint working and backup paths, where the wavelength channels on the backup path are dedicated for that lightpath. Mesh-restored lightpaths are protected using diversely-routed SRG-disjoint working and backup lightpaths, where the wavelength channels on the backup route are shared with other mesh-restored lightpaths

### E. Lightpath Databases

Each optical switch has a lightpath database that contains an entry for each lightpath that traverses the switch (on its working or backup routes). The lightpath database also contains entries for lightpaths that are dropped at that switch. In particular, each lightpath database entry contains the following fields:

- lightpath ID, and attributes;
- lightpath type: unprotected, 1 + 1 (node/SRG) diverse or mesh-restored;
- lightpath route (working/backup).

The lightpath database is updated when a lightpath is provisioned, and when its attributes change. It is noted that the lightpath database at an optical switch does not contain paths that do not traverse that switch. The reason for this is that updating information regarding all lightpaths in the optical network at each optical switch may not be scalable for a dynamically provisioned optical network.

Derived from the lightpath database is a sharing-database that contains, for each switch  $N$ , and for each SRG adjacent to  $N$ , and for each shared channel on that SRG, the lightpaths whose backup routes are routed over that channel. When a new lightpath is provisioned (deprovisioned), it is added (removed) to the sharing database at each switch on its backup route.

The network management system contains the snapshot of the entire network. In particular, it contains the topology database, lightpath databases at each optical switch, and the sharing database at each optical switch.

### F. Lightpath Setup

After a lightpath request is initiated, the procedure to set up the lightpath consists of two tasks:

- 1) routing;
- 2) signaling.

Routing involves computing the exact path from the ingress termination point to the egress termination point across the mesh optical network. For protected lightpaths, an alternate route for the backup lightpath also needs to be determined. Signaling [7] deals with coordinating and setting up the cross-connects at the ingress, egress, and the intermediate optical switches along the route of the lightpath. For protected lightpaths, signaling also needs to be performed for the backup lightpath. It is assumed that explicit routing is used for constructing lightpath routes. Explicit routes may be valuable for traffic engineering and optimizations in the network. The route for the lightpath is determined from the topology database and can be determined by the ingress switch of the lightpath request. Alternatively, the route could be determined by a network management system. With the OSPF extensions described in Section II-D, the ingress optical switch has (for purposes of routing) a representation of the full physical network topology and the available resources on each link.

Upon determination of the route (either at the ingress switch, or by the management system), the ingress switch starts signaling to set up the switches along the lightpath route. Signaling may be performed using extensions to standard signaling protocols such as CR-LDP [5] or RSVP [6]. For unprotected routes, signaling is performed on the working route. For 1 + 1 protected routes, signaling is performed on both the working and backup routes, and switches are configured on both the working and backup routes. For mesh-restored lightpaths, signaling is performed on the working and backup routes, however, switches are configured on the working route, and wavelength channels are soft-reserved on the restoration route by updating the lightpath databases and sharing databases on the backup route (i.e., capacity is reserved, but the optical switches on the backup path are not configured).

For dynamic provisioning, when the time-scales are sufficiently long, mesh-restored lightpaths can be routed from a centralized management system that has access to the network inventory, including the topology database, lightpath databases and the sharing databases. However, when time-scales for dynamic provisioning shrink, then the centralized approach to mesh-restored lightpath routing is infeasible, and then,

mesh-restored lightpaths are assumed to be routed by the ingress optical switch from its local databases.

### G. Lightpath Restoration

For a  $1 + 1$  protected lightpath, nodes adjacent to the failed channel upon detecting the channel failure, send failure messages to the end-nodes of the failed lightpath. The failure message is relayed by intermediate nodes on the working route. Upon receiving the failure message, each end-node switches to the backup lightpath. The restoration time for a  $1 + 1$  protected lightpath is approximated by the sum of the propagation time of optical signals on the working route, and the switching time of the end-nodes.

For a mesh-restored lightpath, nodes adjacent to the failed channel, upon detecting the channel failure, send failure messages to the end-nodes of the failed lightpath. The failure message is relayed by intermediate nodes on the working route. The end nodes, upon receiving the failure message, start signaling on the pre-computed backup path by sending a switch set-up message. Intermediate nodes on the backup path establish the appropriate switch setting, and relay the switch set-up message. The end-nodes upon receiving the switch set-up message switch to the backup lightpath. The restoration time for a mesh-restored lightpath is dominated by the message processing times on the backup path, which is proportional to the number of hops on the backup path.

## III. TOPOLOGY MODELING AND ROUTING

Each optical switch has a representation of the full physical network topology and the available wavelength resources. These are obtained and updated via OSPF opaque LSAs as described in Section II-D. Route computation is performed using this topology database by transforming the topology database into a labeled directed graph.

### A. Modeling the Topology as a Labeled Directed Graph

The topology database obtained via OSPF is converted into a Labeled Directed Graph. Vertices and links in this graph correspond to switches and links in the optical network. Labels on the edges of the graph indicate the SRG information, and the cost and resource information, and enforce the actual connectivity of the topology.

Finding routes in the optical network boils down to finding paths in the Labeled Directed Graph. For example, the shortest-cost path between two nodes can be found in the Labeled Directed Graph by using well-known algorithms such as the Bellman–Ford Algorithm [10]. Finding SRG-disjoint routes in the optical network boils down to finding link/node disjoint routes in the Labeled Directed Graph. This can be accomplished, for example, by utilizing well-known algorithms such as the Suurballe’s algorithm [8], [9].

### B. ROUTING

1) *Unprotected Lightpaths*: An unprotected lightpath from optical switch  $A$  to optical switch  $Z$  is a path in the network wherein each link of the path has dedicated capacity allocated to the lightpath (and is used to carry traffic). A meaningful

notion of the cost for an unprotected lightpath is the sum of the costs of the links of the lightpath. A routing algorithm can therefore pick the minimum cost route while selecting a route among all routes from  $A$  to  $Z$ . A modified version of the Bellman–Ford algorithm [2] may be used to compute the route for an unprotected lightpath. In this case, the routing algorithm constructs the Labeled Directed Graph from the topology database as described in Section III-A. Then it runs a modified version of Bellman–Ford algorithm to compute the shortest-cost route valid route in this Labeled Directed Graph. The modified Bellman–Ford algorithm takes into account the edge labels in the Labeled Directed Graph, while determining the route.

2) *1 + 1 Protected Lightpaths*: A  $1 + 1$  protected lightpath from node  $A$  to node  $Z$  is a pair of node (or SRG) diverse paths in the network where one of the routes is a working route, and the other is the protection route. Each link of both the working and protection routes has dedicated capacity allocated to the lightpath. A meaningful notion of the cost for a  $1 + 1$  lightpath is the sum of the costs of the links of the working lightpath and the protection lightpath. A routing algorithm can therefore pick the minimum combined cost node (or SRG) diverse pair of routes from  $A$  to  $Z$  while selecting among all such node (or SRG) disjoint route pairs from  $A$  to  $Z$ .

The routing algorithm for  $1 + 1$  protected lightpaths first constructs the Labeled Directed Graph from the topology database as described in Section III-A. Then it may use one of several algorithms for finding node (or SRG) disjoint routes such as the Suurballe’s algorithm [8], [9] in the Labeled Directed Graph. The version of Suurballe’s algorithm in turn may use the modified version of the Bellman–Ford algorithm to compute the routes in the Labeled Directed Graph. The modified version of the Suurballe’s algorithm takes into account the edge labels in the Labeled Directed Graph while determining the route in the Labeled Directed Graph.

The above approach for routing  $1 + 1$  protected lightpaths suggests a simple heuristic approach for routing mesh restored lightpaths called “ $1 + 1$  with sharing.” In  $1 + 1$  with sharing, a pair of  $1 + 1$  SRG-disjoint routes is determined as described above, and then one of the routes is chosen as the primary path, and the other route is chosen as the backup path, and channels are shared on as many SRGs as possible on the backup path.

3) *Mesh Restored Lightpaths*: A mesh restored lightpath from node  $A$  to node  $Z$  is a pair of SRG-disjoint paths in the network where one of the routes is a working route, and the other is a backup route. Each link of the working route has dedicated capacity allocated to the lightpath (and carries traffic under normal conditions). Each link of the backup route also has capacity dedicated to this lightpath; however, the capacity dedicated on the backup path can be shared with backup paths for other mesh-restored lightpaths. For any two mesh restored lightpaths  $p1$  (working route =  $w1$ , backup route =  $b1$ ) and  $p2$  (working route =  $w2$ , backup route =  $b2$ ),  $b1$  and  $b2$  can share an SRG if  $w1$  and  $w2$  are SRG disjoint. The above sharing condition ensures that all mesh-restored lightpaths can be restored after any single SRG failure. When any single SRG fails, all mesh restored lightpaths whose working route include the SRG, can be routed on their backup routes without any contention for capacity.

Each switch maintains a sharing database as described in Section II-E, which enables a routing approach to decide if a given lightpath can share wavelength channels with other existing lightpaths. For a given shared wavelength channel  $c$  at a node, a lightpath with working route  $w$  can share  $c$  with other mesh-restored lightpaths if  $w$  is SRG-disjoint with the working routes of all lightpaths that currently use  $c$  on their backup paths. A meaningful notion of the cost of a mesh restored lightpath  $p$  whose working route is  $w$  and backup route is  $b$  is defined as the sum of the costs of the links in  $w$  and costs of the un-shared links in  $b$ . This definition sums up the costs of “new” capacity allocated to the lightpath. Additionally, it may be a requirement that the number of hops on the backup path be bounded. This constraint avoids long backup routes that will violate the restoration time guarantee. A routing algorithm for mesh restored lightpaths can therefore pick the minimum-cost SRG-disjoint route pair from  $A$  to  $Z$  while selecting among all SRG-disjoint route pairs from  $A$  to  $Z$ . The hop-slack parameter indicates the number of additional hops allowed for a backup path over the shortest-hop alternate path that is SRG disjoint from the working route. As the hop-slack parameter increases, longer backup paths that can potentially share capacity better are allowed.

The *mesh restored lightpath routing* Problem is defined as follows: Given the topology database and sharing table at each optical switch, and given switch  $A$  and switch  $Z$ , and a required Type (OC-48, OC-192) of a lightpath, the problem is to find a *feasible* SRG-disjoint working and backup routes from  $A$  to  $Z$ , such that each hop of the working route has available capacity, and each hop on the backup route has either shared capacity or dedicated capacity. To bound the restoration time, an additional constraint may be placed on the number of hops of the backup path. Even without the hop constraint on the backup path, the Mesh Restored Lightpath feasibility problem is NP-complete [11]. The above formulation of the problem merely seeks a *feasible* SRG-disjoint working and backup routes. It is noted that it is not required to find the minimum cost SRG-disjoint pair among all SRG-disjoint working and shared backup routes between  $A$  and  $Z$ . The reason for this is that even the feasibility problem (without the hop distance constraint on the backup route) is NP-complete. The intuitive reason why this problem is NP-complete is that, in the worst case, an algorithm can be forced to enumerate all paths from  $A$  to  $Z$ . There are several heuristic approaches to solving the above routing problem. These approaches rely on simple modifications to standard routing algorithms such as the Suurballe’s algorithm [8], [9] and the Bellman–Ford algorithm [10] that are well-known in the literature. Such heuristics have a running time that is proportional to the square of the number of nodes, and therefore are scalable to large networks.

#### IV. PERFORMANCE RESULTS

We performed simulations of dynamic provisioning on several representative backbone mesh topologies. The results in this paper are reported for the network topology illustrated in Fig. 4 that has 17 nodes and 24 links with an average node degree of




Fig. 4. Representative mesh optical network topology for performance studies.

2.8. Results reported for this topology are representative of results for other backbone mesh network topologies with similar average node degree. For performance studies, we only consider default SRGs, and assume that each wavelength channel has a default cost of 1 unit. The lightpath requests are selected such that a lightpath is equally likely to have any node-pair as its source and destination nodes. We consider a load parameter  $k$ , which indicates the multiple of a full-mesh demand of  $n * (n - 1) / 2$  lightpaths, where  $n$  is the number of nodes. Backbone demand between node-pairs is not expected to be uniformly distributed, we expect some skew in the demand with some hub nodes sourcing and sinking more lightpaths than other nodes. However, we expect that the results obtained with the assumption of uniformly distributed demand will be robust in the sense that they will continue to be representative of any dense traffic demand.

First, we examine the capacity requirements (in terms of the total number of wavelength channels to route the lightpath demand) for each of the lightpath types for different values of the load parameter ranging from 0.2 to 4.0. Fig. 5 illustrates the capacity requirement for 1 + 1 protected node diverse, 1 + 1 protected link-diverse, 1 + 1 with sharing, and mesh-restored lightpaths for different values of the load parameter. The capacity requirement in each case is normalized with respect to the capacity requirement for unprotected lightpaths. We observe that the total capacity requirement for 1 + 1 protected node-diverse lightpaths is around 270–275% regardless of the load. The total capacity requirement for 1 + 1 protected link-diverse lightpaths is around 260–265% regardless of the load. Node-diverse 1 + 1 protected lightpaths require around 5% more protection capacity than 1 + 1 protected link-diverse lightpaths. This extra capacity is the cost of providing a stronger protection guarantee of protecting against all single-node failures in addition to protecting against all single-link failures. The capacity requirement for 1 + 1 with sharing decreases from 195% at a load of 0.2 to 186% at a load of 0.7. Beyond a load of 0.7 the decrease in capacity requirement is marginal. For mesh-restored lightpaths, the capacity requirements decreases from 185% at a load of 0.2 to 170% at a load of 0.7. Beyond a load of 0.7, the capacity requirement remains around 165%. Gains in capacity requirements are achieved due to sharing of wavelength channels on the backup path. As the number of lightpaths in the network




Fig. 5. Total capacity requirement versus load. Capacity requirement is normalized with respect to the unprotected working capacity.




Fig. 6. Sequence of topologies with increasing average node degree constructed by adding chords to a ring.




Fig. 7. Protection capacity requirement against the average node degree for mesh-restored lightpaths.

increases, there is more opportunity for sharing of lightpaths. However, the sharing gains saturate as load increases.

We then examine the sensitivity of the benefits of sharing of wavelength channels on the backup path to the average node degree of the topology. For this, we examine a sequence of topologies with increasing node degree constructed as follows: we start with a ring where each node has degree 2 (and is the smallest topology that admits diverse routes between all node pairs). Then we add chords to the ring to achieve the desired

connectivity in terms of node-degree, as illustrated in Fig. 6. Such topologies are not very practical in reality; however, they illustrate the dependence of sharing behavior on the average node degree, and we expect such behavior to manifest in realistic topologies as well.

Fig. 7 plots the protection capacity (normalized as a percentage of the working capacity) requirement versus node degree for a 18 node network when the degree of the topology varies from 2 to 4 (i.e., we go from a ring to denser topologies). For a ring-network

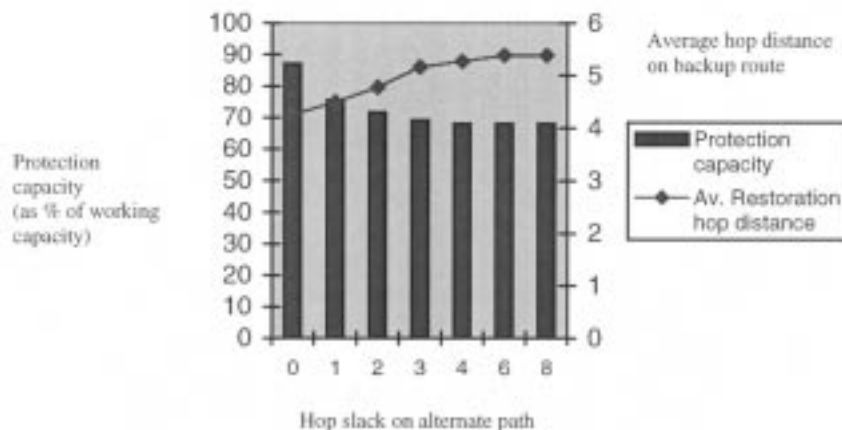


Fig. 8. Protection capacity requirement, and average backup path hop distance versus hop slack on backup path.

with node degree 2, and for a full-mesh demand, it can be easily shown that the shared protection capacity is exactly equal to the working capacity. In Fig. 7, for the 18 node ring, the protection capacity is about 103% of working capacity due to the fact that the demand is random uniformly distributed, and deviates slightly from a full-mesh demand. As the average node-degree increases from 2, the protection capacity drops rapidly, and when the average node degree is 2.5, the protection capacity is about 50% of working capacity. At an average node degree of 3, the protection capacity is about 45% of working capacity. Sharing benefits increases rapidly at first with increasing node degree. However, with increasing degree, the marginal improvement in sharing gains decreases rapidly. Interestingly, we observe significant sharing gains even when the network is quite sparse (i.e., at an average node degree of 2.5).

We then examine the effect of the hop distance of the backup path on the sharing capacity gains. We vary the hop distance on the backup path by varying the hop-slack parameter in the mesh-restoration routing algorithm. The hop-slack parameter indicates the number of additional hops allowed for a backup path over the shortest-hop alternate path. As the hop-slack parameter increases, longer backup paths that can potentially share capacity better are allowed. Therefore we expect better sharing when we allow longer backup paths. However, longer backup paths lead to longer restoration time for restoring lightpaths since the restoration time is proportional to the hop-distance of the backup path. Fig. 8 illustrates the protection capacity requirement and the average backup path hop-distance versus the hop-slack allowed on the backup path. We observe that as the hop-slack is increased from 0 to 3 (i.e., by allowing backup paths up to at most 3 hops longer than the shortest backup path) there is a reduction in protection capacity of almost 17%. However, beyond a hop-slack of 3 hops, the gain in protection capacity is marginal.

## V. CONCLUSION

In this paper, we described the architecture for the dynamic provisioning of lightpaths in an optical network and outline approaches for routing unprotected, 1 + 1 protected and mesh-restored lightpaths.

We empirically analyze via simulations the capacity performance of the dynamically provisioning approaches on representative mesh network topologies. The analysis of the efficiency of network utilization of dynamic provisioning focuses on the spare capacity needed for protection, and in particular on the impact of sharing of wavelength channels for mesh-restored lightpaths. The main conclusions from the performance studies are:

- 1) significant sharing capacity gains are achieved for mesh-restored lightpaths with dynamic provisioning even at moderate loads;
- 2) significant sharing capacity gains are achieved with dynamic provisioning even for sparse topologies; and
- 3) the average hop-distance on the backup path (which is proportional to the restoration time) can be traded-off against the required protection capacity.

## REFERENCES

- [1] T. E. Stern and K. Bala, *Multiwavelength Optical Networks: A Layered Approach*. Reading, MA: Addison-Wesley, 1999.
- [2] S. Chaudhuri, G. Hjalmtysson, and J. Yates, "Control of lightpaths in an optical network," in *Optical Internetworking Forum*, Jan. 2000, OIF2000.004.0.
- [3] J. Moy, "OSPF Version 2," RFC 1247, July 1991.
- [4] K. Bala, S. Giebl, R. Ramamurthy, W. Russ, and B. Tang, "IP centric control and signaling for optical lightpaths," in *Optical Internetworking Forum*, Jan. 2000, OIF2000.019.
- [5] B. Tang, D. Saha, and B. Rajagopalan, "Extensions to CR-LDP for path establishment in optical networks," Internet Draft (Work in progress), draft-tang-crl-dp-optical-00.txt, March 2000, to be published.
- [6] D. Saha, B. Rajagopalan, and B. Tang, "RSVP extensions for signaling optical paths," Internet Draft (Work in progress), draft-saha-rsvp-optical-signaling-00.txt, Mar. 2000, to be published.
- [7] B. Rajagopalan, D. Saha, B. Tang, and K. Bala, "Signaling for automated establishment and restoration of paths in optical mesh networks," Internet Draft (Work in progress), draft-rstb-optical-signaling-00.txt, Mar. 2000, to be published.
- [8] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Norwell, MA: Kluwer, 1998.
- [9] J. W. Suurballe, "Disjoint paths in a network," *Networks*, vol. 4, 1974.
- [10] T. Cormen, C. Leiserson, and R. Rivest, *Introduction to Algorithms*. New York: McGraw-Hill, 1990.
- [11] M. R. Garey and D. S. Johnson, "Computers and intractability," in *A Guide to the Theory of NP-Completeness*. San Francisco, CA: Freeman, 1979.

**Ramu Ramamuthy** received the B.Tech. degree from the Indian Institute of Technology, Madras, and M.S. and Ph.D. degrees from the University of California, Davis.

He is currently at Tellium, Inc., where he works on the design of algorithms and protocols for dynamic provisioning and restoration in optical networks. Prior to this, he was at Telecordia Technologies, where he worked on network control and management of IP/WDM optical networks.

**Zbigniew Bogdanowicz** received Masters degree in electrical engineering from Wroclaw Polytechnic, Poland, and the Ph.D. degree in computer science from Stevens Institute of Technology, Hoboken, NJ.

He is currently at Tellium, Inc., where he leads the Planning and Simulation Tools Group. Prior to joining Tellium, he worked for 15 years as Member and Distinguished Member of Technical Staff in Bell Laboratories. His main research interests include Graph Theory, Combinatorial Optimization, Routing Protocols, and Network Design Algorithms.

**Shahrokh Samieian** currently works at Tellium, Inc., as a Senior Network Architect in the optical networks design tools group responsible for design and implementation of planning and simulation tools. Prior to joining Tellium, he worked at Nortel Networks as a Senior Member of Scientific Staff in the optical network evolution group responsible for forward looking optical product planning as well as metro and long-haul network planning for major service providers such as MCI/WorldCom and Qwest.

**Debanjan Saha** Debanjan Saha received M.S. and Ph.D. degrees from the University of Maryland at College Park, and B.Tech. degree from Indian Institute of Technology, India, all in computer science.

He is currently with Tellium, Inc. Previously, he spent several years at IBM Research and Lucent Bell Labs. He is actively involved with various standards bodies, serves as editor of international journals and magazines, and on technical committees of national and international conferences. He has authored numerous technical articles on various networking topics, and is a frequent speaker at academic and industrial events.

**Bala Rajagopalan** received the B. Tech degree from the Indian Institute of Technology, and the M.S and Ph.D degrees from the University of Illinois at Urbana-Champaign.

He is currently with Tellium, Inc. He is responsible for Tellium's IP over optical and optical internetworking architectures. Prior to joining Tellium, he was with AT&T Bell Laboratories, Bellcore, and NEC C&C Research Laboratories, where he developed protocols and algorithms for wireless data networking, Internet routing and multicast, quality-of-service (QoS), traffic engineering, and IP switching. He has published numerous papers and delivered many invited lectures on topics covering his areas of interest. He is also very active in standardization activities and he has made significant contributions to the IETF, the ATM Forum and the Optical Interworking Forum (OIF).

**Sudipta Sengupta** received the M.S. degree from Massachusetts Institute of Technology, Cambridge, and the B.Tech. degree from Indian Institute of Technology, Kanpur, India, both in computer science.

He is currently with Tellium, Inc., where he works on the design of algorithms and protocols for dynamic provisioning and restoration in optical networks. Prior to this, he was with Oracle, where he was involved in the design and development of the wireless networking platform for Oracle Mobile Applications. He was also at Bell Labs Research, Lucent Technologies, where he worked on online QoS routing and traffic engineering under the MPLS framework.

Dr. Sengupta received the President of India Gold Medal at IIT-Kanpur for academic excellence.

**Sid Chaudhuri** (M'97) received the Ph.D. degree in physics from the University of Pittsburgh, PA, in 1981.

He is currently with Tellium, Inc., where he leads the network architecture department to develop mesh-networking protocols for restoration and dynamic provisioning, and network planning and simulation tools. Prior to Tellium, he was with AT&T Bell Labs and AT&T Labs for 15 years. While at AT&T, he led the development of AT&T's Core Transport Network architecture. He has published more than 30 research papers on electronic and optical properties of impurities in bulk semiconductors, semiconductor quantum well structures and optical networking. He also serves as the Vice President of the Board of Directors in the Optical Internetworking Forum (OIF).

**Krishna Bala** received the Ph.D. degree from Columbia University, NY, for his work on routing algorithms for optical networks.

He is a Founder and Chief Technical Officer of Tellium, Inc. He is a lead system architect for the Aurora Optical Switch product. He spent five years with Bellcore as Lead Architect for the multiwavelength optical networks project for which he received the Bellcore President's Award of Excellence. He has several patents and publications in the area of optical switching and is a coauthor of a book on the subject.